

REMARKS

This Amendment is filed in response to the Office Action mailed on February 18, 2010. All objections and rejections are respectfully traversed.

Claims 1-8, 12, 14-25, 29, 31-35, 37-39, 41-42, 46-50 are in the case.

No new claims have been added.

Claims 1, 5-7, 14, 18, 20, 22, 24, 31-34, and 38-39 have been amended.

Interview Summary

Applicant would like to thank Examiner Lovel for conducting the Applicant Initiated Interview on April 15, 2010 and for helping to advance this Application closer to allowance. Generally, as will be elaborated upon in greater detail below, the issue discussed involved Applicant's use of **scanning a root of each PCPI comprising a (particular) directory tree to create an index associated with** (that particular) *directory tree*. Specifically, Applicant discussed that even though Mani teaches creating a database of various catalogues, each catalogue is specific to each single snapshot (and not specific to the data of the snapshot), which means that Mani teaches creating a non-specific database of snapshots. In contrast, Applicant claims **scanning a root of each PCPI comprising the (particular) directory tree to create an index associated with** (that particular) *directory tree*. As such, because Mani teaches cataloguing all snapshots because they are snapshots, as opposed to cataloguing particular snapshots because those particular snapshots have been determined by a scan of the root of those particular snapshot to comprise particular data, Mani fails to teach or suggest Applicant's claimed novel and non-obvious **scanning a root of each PCPI comprising the (particular) directory tree to create an index associated with** (that particular) *directory tree*.

While no agreement was officially reached, Examiner noted that a closer look at the prior art would be required to verify Applicant's contentions and that another search may be conducted. If a new search results in new art, Applicant respectfully requests that Examiner contact the undersigned attorney to discuss the art before issuing the next Of-

fice Action. Examiner is encouraged to contact the undersigned attorney with any questions.

Claim Objections

At paragraph 5 of the Office Action, Examiner objected to claims 5-7, 14-15, 18, 22, 27, 31-32, 34, and 38-39 due to informalities. Applicant thanks Examiner for noting such informalities. Claims 5-7, 14, 18, 22, 27, 31-32, 34, and 38-39 have been amended and are believed to be in condition for allowance. Applicant respectfully contends that claim 15 does not require the amendment suggested by Examiner.

35 U.S.C. §101-Clarifications

In response to Examiner’s statements at paragraphs 6 and 7 of the Office Action, Applicant respectfully notes that all currently pending claims are directed only to statutory subject matter that meet the requirements under 35 U.S.C. §101.

Rejections Under 35 U.S.C. §103

At paragraph 10 of the Office Action, claims 1-2, 6, 15-19, 31-35, 38, and 49-50 were rejected under 35 U.S.C. §103(a) as being obvious over Mani-Meitav et al., U.S. Patent Application Publication No. 2005/0216788 published on September 29, 2005 (hereinafter “Mani”), in view of Dewey et al., U.S. Patent No. 7,529,778 issued on May 5, 2009 (hereinafter “Dewey”).

Applicant’s claimed novel and non-obvious invention, as set forth in representative claim 1, comprises in part:

The present invention, as set forth in representative claim 1, comprises in part:

1. A system for indexing and manipulating backup data stored on a destination storage system, comprising:
 - one or more source servers configured to transmit the backup data to the destination storage system;
 - a management application executed by a processor, wherein the **management application** is **configured to** (a) communicate with the

destination storage system and further configured to access data identifiers related to the backup data organized in a directory tree structure representing a plurality of persistent consistency point images (PCPIs) of the backup data, wherein each PCPI is associated with a creation time, (b) ***scan a root of each PCPI comprising the directory tree to generate an index*** of directories, files, or qtrees ***associated with the directory tree***, and (c) organize the data identifiers into a structure that enables the backup data to be displayed; and

a user interface to select a directory, file, or qtree to view, wherein the management application is further configured to return a list of the selected directory, file, or qtree and one or more versions of the selected directory, file, or qtree.

Mani teaches, in relevant part, cataloguing snapshots so that in recovery mode, “a chosen version of a certain set of data may be selected” [0138]. Notably, each catalog is associated with a specific snapshot, and is used to find data related to one snapshot [0166]. Put another way, each single snapshot is associated with a catalogue specific to the snapshot. A database of the catalogues contains the various catalogues, each catalogue which is specific to each single snapshot. Mani explicitly states that the database is generated by:

...file-level recovery procedure enabling the restoration of data to file-level data configuration, by reading and retrieving data stored in the repository 15 in the block-level data configuration...The analysis process scans the raw block-level backup image in the repository 15 for file-system organization, by first reading the type of file-system from the image.

Dewey teaches, in relevant part, obtaining a list of earlier file versions that exist on a temporal shadow volume (col. 8, lines 60-65). Specifically, to obtain such a list, Dewey explicitly states that locating the file versions is accomplished by first obtaining at the local (client) machine a list of the available shadow volumes (col. 9, lines 23-28; *see also* Fig. 3). If any shadow volume exists, this information is returned in the shadow volume list (col. 9, lines 49-51).

Applicant respectfully urges that Mani, taken singly or in any combination with Dewey, does not disclose, teach, or suggest Applicant's claimed novel and non-obvious **management application configured to scan a root of each PCPI comprising the directory tree to generate an index associated with the directory tree.**

Applicant claims, in part, a **management application configured to** communicate with a destination storage system and further configured to access data identifiers related to the backup data organized in a directory tree structure representing a plurality of persistent consistency point images (PCPIs) of the backup data. Each PCPI may pertain to a particular directory tree or, e.g., file. In other words, if numerous PCPIs are taken of a storage system containing that particular directory tree, then each of those PCPIs may also comprise that particular directory tree. Notably, as more PCPIs of the particular directory tree are created at different points-in-time, each of those multiple PCPIs also comprise that directory tree.

However, while some PCPIs may contain that particular directory tree, not every PCPI necessarily does. It may be desirable to generate an index of those PCPIs associated with a particular directory tree spanning multiple points-in-time (i.e., **each PCPI comprising the directory tree**), rather than generating an index of every PCPI making up a single point-in-time image of the entire active file system. In other words, where information on a particular directory is requested, the full range of those PCPIs containing that particular directory (e.g. the root of each relevant PCPI comprising that particular directory tree) may be located and read to generate an index associated with that particular directory tree. Applicant illustratively accomplishes this, in part, by using a **management application configured to scan a root of each PCPI comprising the (particular) directory tree to create an index associated with** (that particular) **directory tree**.

Applicant respectfully contends that Mani fails to teach or suggest Applicant's claimed novel and non-obvious scanning a root of each PCPI comprising a (particular) directory tree to **create an index associated with** (that particular) **directory tree**. Specifi-

cally, Mani explicitly states that each catalog is associated with a specific snapshot, and is used to find data related to one snapshot. In other words, only a single snapshot is to be examined as each catalogue only pertains to data related to one snapshot. Thus, even though Mani teaches creating a database of various catalogues, each catalogue is specific to each single snapshot (and not specific to the data of the snapshot), which means that Mani teaches creating a non-specific database of snapshots. In contrast, Applicant claims scanning a root of each PCPI comprising the (particular) directory tree to **create an index associated with** (that particular) *directory tree*. As such, because Mani teaches cataloguing all snapshots because they are snapshots, as opposed to cataloguing particular snapshots because those particular snapshots have been determined by a scan of the root of those particular snapshot to comprise particular data, Mani fails to teach or suggest Applicant's claimed novel and non-obvious scanning a root of each PCPI comprising the (particular) directory tree to **create an index associated with** (that particular) *directory tree*.

Furthermore, Applicant respectfully contends that Mani also fails to teach or suggest Applicant's claimed novel and non-obvious **management application configured to scan a root of each PCPI comprising the directory tree to generate an index associated with the directory tree**. Specifically, while Mani teaches cataloguing "snapshots", Mani teaches that the cataloguing is generated by reading and retrieving data stored in a repository in the block-level data configuration and scanning the raw block-level backup image in the repository by reading the type of file-system from the image. Notably, neither of those steps is taught or suggested to be the same as scanning a root of each PCPI. In contrast, Applicant claims a management application configured to **scan a root of each PCPI comprising the directory tree** to generate an index associated with the directory tree. As such, because Mani catalogues "snapshots" by methods other than scanning a root of each PCPI comprising the (requested) directory tree, Mani fails to teach or suggest Applicant's claimed novel and non-obvious **management application configured**

to scan a root of each PCPI comprising the directory tree to generate an index associated with the directory tree.

Applicant respectfully contends that Dewey fails to teach or suggest Applicant's claimed novel and non-obvious **management application configured to scan a root of each PCPI comprising the directory tree to generate an index associated with the directory tree**. Specifically, while Dewey teaches obtaining a list of earlier file versions, Dewey also states that such a list is generated by obtaining at the local (client) machine a list of all the available shadow volumes. In other words, if any shadow volume exists, this information is returned in the shadow volume list. Notably, returning a list of all shadow volumes is not taught or suggested by Dewey to be the same as scanning a root of each PCPI, nor is such an assertion made by Examiner. In contrast, Applicant claims a management application configured to **scan a root of each PCPI comprising the directory tree** to generate an index associated with the directory tree. As such, because Dewey teaches generating a list of files by methods other than scanning a root of each PCPI comprising the (requested) directory tree, Dewey fails to teach or suggest Applicant's claimed novel and non-obvious **management application configured to scan a root of each PCPI comprising the directory tree to generate an index associated with the directory tree**.

Based on the foregoing, Applicant respectfully urges that Mani, taken singly or in any combination with Dewey, fails to render the presently claimed invention obvious under 35 U.S.C. §103(a). Mani and Dewey, taken singly or in any combination, do not disclose, teach, or suggest Applicant's claimed novel and non-obvious **management application configured to scan a root of each PCPI comprising the directory tree to generate an index associated with the directory tree**.

At paragraph 11 of the Office Action, claims 3-5, 20-23, 37, and 41 were rejected under 35 U.S.C. §103(a) as being obvious over Mani, in view of Dewey, and in further view of Armangau, U.S. Patent No. 6,434,681 (hereinafter “Armangau”).

Applicant respectfully notes that claims 3-5, 20-23, 37, and 41 are dependent claims that depend from independent claims believed to be in condition for allowance. Accordingly, claims 3-5, 20-23, 37, and 41 are believed to be in condition for allowance.

At paragraph 12 of the Office Action, claims 7-8, 12, 14, 24-25, and 29 were rejected under 35 U.S.C. §103(a) as being obvious over Mani, in view of Dewey, and in further view of Arakawa et al., U.S. Patent Application Publication No. 2003/0131207 (hereinafter “Arakawa”).

Applicant respectfully notes that claims 7-8, 12, 14, 24-25, and 29 are dependent claims that depend from independent claims believed to be in condition for allowance. Accordingly, claims 7-8, 12, 14, 24-25, and 29 are believed to be in condition for allowance.

At paragraph 13 of the Office Action, claims 42 and 46-48 were rejected under 35 U.S.C. §103(a) as being obvious over Mani, in view of Dewey, and in further view of Arakawa.

The present invention, as set forth in representative claim 42, comprises in part:

42. A system, comprising:

a source storage system configured to generate a plurality of persistent consistency point images (PCPIs) associated with a particular directory tree, and further configured to transfer the plurality of PCPIs to a destination storage system;

the destination storage system configured to execute a management client, wherein the management client is configured to organize the plurality of PCPIs into an index using a database to allow the plurality of PCPIs to be displayed in (a) a listing of source data entries indexed by the

particular directory tree, wherein each PCPI of the particular directory tree is created at one or more different times (b) a listing of source data entries indexed by names of the source storage system, and (c) a listing of source data entries indexed by names of volumes of the destination storage system in which backup data from the source storage system resides; and

an interface configured to select a data entry for the particular directory tree, and the **management client** further **configured to return a list of the plurality of PCPIs associated with the particular directory tree**.

As noted above, Mani teaches that each catalog is associated with a specific snapshot, and is used to find data related to one snapshot. As such, because Mani teaches finding data related to only one snapshot, Mani fails to teach or suggest Applicant's claimed novel and non-obvious **management client configured to return a list of the plurality of PCPIs associated with the particular directory tree**.

As also noted above, while Dewey teaches obtaining a list of earlier file versions, Dewey also states that such a list is generated by obtaining at the local (client) machine a list of all the available shadow volumes. In other words, if any shadow volume exists, this information is returned in the shadow volume list. Notably, returning a list of all shadow volumes is not taught or suggested by Dewey to be the same as **returning a list of the plurality of PCPIs associated with the particular directory tree**. As such, because Dewey does not teach returning a list of the PCPIs associated with the requested file, Dewey fails to teach or suggest Applicant's claimed novel and non-obvious **management client configured to return a list of the plurality of PCPIs associated with the particular directory tree**.

Arakawa teaches a virtualized volume snapshot formation method. However, Arakawa fails to teach or suggest Applicant's claimed novel and non-obvious **management client configured to return a list of the plurality of PCPIs associated with the particular directory tree**.

Based on the foregoing, Applicant respectfully urges that Mani, taken singly or in any combination with Dewey and/or Arakawa, fails to render the presently claimed invention obvious under 35 U.S.C. §103(a). Mani and/or Dewey and/or Arakawa, taken singly or in any combination, do not disclose, teach, or suggest Applicant's claimed novel and non-obvious **management client configured to return a list of the plurality of PCPIs associated with the particular directory tree.**

Applicant's Interpretation of the Prior Art

Applicant's interpretation of the prior art was derived, in part, from the following excerpts:

Mani

[0138] Each single snapshot is associated with a catalogue specific to the snapshot and listing attributes pertaining to the snapshot, such as snapshot version, number, device from which the snapshot was read, source data location, date, and time of the snapshot. Catalogued data permits to retrace the chronological history of any portion of backed up data, and may be used for calling up a string of successive snapshots, out of storage from a data storage repository 15. A database of catalogues contains the various catalogues, which are specific to each single snapshot. Hence, in recovery mode, a chosen version of a certain set of data may be selected, then "stringed up" and restored, starting from say one given snapshot-point in time to another snapshot-point in time. (emphasis added)

[0166] The FBSRD operates a file-level recovery procedure enabling the restoration of data to file-level data configuration, by reading and retrieving data stored in the repository 15 in the block-level data configuration... The analysis process scans the raw block-level backup image in the repository 15 for file-system organization, by first reading the type of file-system from the image, and then proceeds by applying a specific analysis method for each type of file-system. The analysis process proceeds iteratively through all elements of the file-system structure until all files and directories have been analyzed. The derived file-system structure is stored in the repository 15 as a catalog database for subsequent use during file-level restoration, when in recovery mode. Each catalog is associated with a specific snapshot, and is used to find data related to one snapshot. (emphasis added)

Dewey

...locating the file versions is accomplished by first obtaining at the local (client) machine a list 300 of the available shadow volumes. To this end, the version locator API 208 includes or is otherwise associated with a client-side shadow volume location mechanism 304, which communicates with a counterpart server-side volume location mechanism 320 to obtain the list 300 of the temporal shadow volumes. (emphasis added)

Then, the shadow volume manager 322 is asked to return information (e.g. a timestamp) for the identified base volume (e.g., 228) for each shadow volume that has been captured therefor. If any shadow volumes exist, this information is then returned in the shadow volume list 300. (col. 9, lines 49-51) (emphasis added)

Conclusion

All new claims and/or claim amendments are believed to be fully supported by Applicant's specification.

All independent claims are believed to be in condition for allowance.

All dependent claims are believed to be dependent from allowable independent claims, and therefore in condition for allowance.

Favorable action is respectfully solicited.

Please charge any additional fee occasioned by this paper to our Deposit Account No. 03-1237.

Respectfully submitted,

/Michael T. Abramson/
Michael T. Abramson
Reg. No. 60,320
CESARI AND MCKENNA, LLP
88 BLACK FALCON AVENUE
BOSTON, MA 02210
Telephone: (617) 951-2500
Facsimile: (617) 951-3927